

Resource Model Equation (7) represents just one exemplary embodiment of the disclosed methods which may be employed to model and balance utilization of I/O capacity and available buffer space for an information management system having multiple storage devices or groups of storage devices. As with Resource Model Equation (4), it will be understood that I/O capacity and available buffer space may be balanced using any other equation, algorithm or other relationship suitable for estimation of I/O capacity and/or buffer space using monitored system I/O performance characteristics and/or other system I/O performance characteristics. For example, the Skew value is just one example of an information management system I/O performance characteristic that may be employed to model or otherwise represent uneven distribution of total I/O demands among a number of multiple storage devices NoD, it being understood that any other information management system I/O performance characteristic suitable for modeling or otherwise representing uneven distribution of total I/O demands among a number of multiple storage devices NoD may be employed.

Resource Modeling for Multiple Storage Devices in Substantially-Unbalanced Workload Environments

In those cases where workload may not be substantially balanced or evenly distributed across multiple storage devices or groups of storage devices (e.g., across disk drives or disk drive groups) a dynamic measurement-based resource model approach may be employed. Although such a measurement-based resource model approach may be implemented under any disk workload distribution (including under substantially balanced workload distributions), it may be most desirable where workload distribution is substantially unbalanced (e.g., in the case of a maximum Skew value of greater than or equal to about 2 for any given one or more storage devices or groups of storage devices). Such a dynamic measurement-based approach becomes more desirable as the possible maximum Skew value increases, due to the difficulty in estimating and pre-configuring a system to handle anticipated, but unknown, future workload distributions.

A dynamic measurement-based resource model may function by actually measuring or monitoring system I/O performance characteristics at least partially reflective of how workload is

distributed across multiple storage devices or groups of devices rather than merely estimating or modeling the distribution. In one embodiment, such a measurement-based resource model may function in conjunction with a storage management processing engine capable of tracking workloads of each storage device or group of storage devices such that the maximal number of 5 viewers to a storage device or group of storage devices, and the maximal aggregated consumption rate to a storage device or group of storage devices may be obtained and considered using the resource model.

One resource model embodiment that may be implemented in substantially unbalanced 10 parallel-functioning multiple storage device environments may employ substantially the same buffer space constraints as described above in relation to single storage device and substantially balanced multiple storage device environments. However, such a resource model may be implemented in a manner that considers additional measured or monitored system I/O 15 performance characteristics to reflect the substantially unbalanced parallel nature of workload distribution in the determination of I/O capacity constraints. These additional monitored performance characteristics may include, but are not limited to, the maximal aggregate consumption rates (“MaxAggRate_perDevice”) that may be expressed as “Max{ $\sum_i \epsilon_{\text{Device}} P_i$; for all devices/groups}”, and the maximal aggregate number of viewers, (“MaxNoV_perDevice”) that may be expressed as “Max{Number of viewers on a device (or a 20 storage device group); for all devices/groups}”.

Once again, cycle time T should be greater than or equal to the maximal aggregate storage device service time for continuous playback, *i.e.*, the maximal aggregate sum of access time and data transfer time for all storage devices or groups of storage devices. Using the above 25 relationships to balance sufficient I/O capacity and sufficient total available buffer space for multiple storage device environments in a manner similar to that employed for Resource Model Equations (4) and (7), range of cycle time T to ensure continuous playback may be defined in one embodiment for substantially unbalanced multiple storage device environments by Resource Model Equation (8A) as follows:

$$\text{MaxNoV_perDevice} * \text{AA} / [1 - \text{MaxAggRate_perDevice} / \text{TR}] \leq T$$

$$\leq B_{\max} / (\sum_{i=1}^{\text{Nov}} P_i) \quad (8A)$$

For an embodiment employing a buffer-sharing scheme, the following equation may be

5 alternatively employed:

$$\text{MaxNoV_perDevice} * \text{AA} / [1 - \text{MaxAggRate_perDevice} / \text{TR}] \leq T$$

$$\leq B_{\max} / [(1 - B_{\text{Save}}) * (\sum_{i=1}^{\text{Nov}} P_i)] \quad (8A')$$

10 In the practice of the disclosed methods and systems, Resource Model Equation (8A) may be employed for I/O admission control and the read-ahead estimation in a manner similar to Resource Model Equations (4) and (7). In one embodiment, the maximal aggregate consumption rates and the maximal aggregate number of viewers may be tracked and utilized in system 100 of FIG. 1 by I/O manager 140 of storage management processing engine 105 to determine whether or not system 100 can support all viewers without compromising quality of playback (e.g., video playback). As with Resource Equations (4) and (7), if values of I/O capacity and buffer space determined from Resource Model Equation (8A) overlap then system 100 cannot support all viewers without compromising quality of playback. However, if a value or range of values for cycle time T exist that will satisfy Resource Model Equation (8A), then system 100 can support 15 all viewers. Assuming the latter to be the case, Resource Model Equation (8A) may be used to determine a range of cycle time T suitable for continuous playback and to give an estimation of optimal read-ahead size for each viewer in a manner similar to that described for Resource Model Equations (4) and (7). Thus, Resource Model Equation (8A) may be employed under unbalanced workload conditions to adjust cycle time T and read-ahead size for existing and new 20 viewers in order to maximize the number of viewers supported by an information management system having multiple storage devices or groups of storage devices.

25